Faster variational inducing input Gaussian process classification

***P. Izmailov*[1], *D. Kropotov*[1],**
izmailovpavel@gmail.com, dmitry.kropotov@gmail.com
[1]Lomonosov Moscow State University, 1 Leninskie Gory, Moscow, Russia

arXiv:1611.06132v1 [cs.LG] 18 Nov 2016

Background: Gaussian processes (GP) provide an elegant and effective approach to learning in kernel machines. This approach leads to a highly interpretable model and allows using the bayesian framework for model adaptation and incorporating the prior knowledge about the problem. GP framework is successfully applied to regression, classification and dimensionality reduction problems. Unfortunately, the standard methods for both GP-regression and GP-classification scale as $\mathcal{O}(n^3)$, where n is the size of the dataset, which makes them inapplicable to big data problems. A variety of methods have been proposed to overcome this limitation both for regression and classification problems. The most successful recent methods are based on the concept of inducing inputs. These methods reduce the computational complexity to $\mathcal{O}(nm^2)$, where m is the number of inducing inputs with m typically much less, than n. In this work we focus on classification. The current state-of-the-art method for this problem is based on stochastic optimization of an evidence lower bound, that depends on $\mathcal{O}(m^2)$ parameters. For complex problems, the required number of inducing points m is fairly big, making the optimization in this method challenging.

Methods: We analyze the structure of variational lower bound that appears in inducing input GP classification. First we notice that using quadratic approximation of several terms in this bound, it is possible to obtain analytical expressions for optimal values of most of the optimization parameters, thus sufficiently reducing the dimension of optimization space. Then we provide two methods for constructing necessary quadratic approximations. One is based on Jaakkola-Jordan bound for logistic function and the other one is derived using Taylor expansion.

Results: We propose two new variational lower bounds for inducing input GP classification that depend on a number of parameters. Then we propose several methods for optimization of these bounds and compare the resulting algorithms with the state-of-the-art approach based on stochastic optimization. Experiments on a bunch of classification datasets show that the new methods perform as well or better than the existing one. However, new methods don't require any tunable parameters and can work in settings within a big range of n and m values thus significantly simplifying training of GP classification models.

Keywords: *Gaussian process; classification; variational inference; big data; inducing inputs; optimization; variational lower bound*

1 Introduction

Gaussian processes (GP) provide a prior over functions and allow finding complex regularities in data. Gaussian processes are successfully used for classification/regression problems and dimensionality reduction [5]. In this work we consider the classification problem only.

Standard methods for GP-classification scale as $\mathcal{O}(n^3)$, where n is the size of a training dataset. This complexity makes them inapplicable to big data problems. Therefore, a variety of methods were introduced to overcome this limitations [8–10]. In the paper we focus on methods based on so called inducing inputs. The paper [6] introduces the inducing inputs approach for training GP models for regression. This approach is based on variational inference and proposes a particular lower bound for marginal likelihood (evidence). This bound is then maximized w.r.t. parameters of kernel function of the Gaussian process, thus fitting the model

to data. The computational complexity of this method is $\mathcal{O}(nm^2)$, where m is the number of inducing inputs used by the model and is assumed to be substantially smaller than n. The paper [1] develops these ideas by showing how to apply stochastic optimization to the evidence lower bound similar to the one used in [6]. However, a new lower bound depends on $\mathcal{O}(m^2)$ variational parameters that makes optimization in the case of big m challenging.

The paper [2] shows how to apply the approach from [1] to the GP-classification problem. It provides a lower bound that can be optimized w.r.t. kernel parameters and variational parameters using stochastic optimization. However, the lower bound derived in [2] is intractable and has to be approximated via Gauss-Hermite quadratures or other integral approximation techniques. This lower bound is also fit for stochastic optimization and depends on $\mathcal{O}(m^2)$ parameters.

In this work we develop a new approach for training inducing input GP models for classification problems. Here we analyze a structure of variational lower bound from [2]. We notice that using quadratic approximation of several terms in this bounds it is possible to obtain analytical expressions for optimal values of the most of optimization parameters thus sufficiently reducing the dimension of optimization space. So, we provide two methods for constructing necessary quadratic approximations — one is based on Jaakkola-Jordan bound for logistic function, and the other one is derived using Taylor expansion.

The paper is organized as follows. In section 2 we describe the standard GP-classification framework and its main limitations. In section 3 we introduce the concept of inducing inputs and derive the evidence lower bound of [2]. Section 4 consists of our main contribution - two new tractable evidence lower bounds and different methods for their optimization. Section 5 provides experimental comparison of our new methods with the existing approach from [2] and the last section concludes the paper.

2 GP-classification model

In this section we review classic Gaussian Processes (GP) framework and its application for classification problems (see [5] for detailed discussion).

2.1 Gaussian process definition

A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution.

We will only consider processes, that take place in a finite-dimensional real space $\mathbb{R}^d$. In this case, f is a Gaussian process, if for any k, for any $\boldsymbol{t}_1, \ldots, \boldsymbol{t}_k \in \mathbb{R}^d$ the joint distribution

$$(f(\boldsymbol{t}_1), \ldots, f(\boldsymbol{t}_k))^\mathsf{T} \sim \mathcal{N}(\boldsymbol{m}_t, \boldsymbol{K}_t)$$

for some $\boldsymbol{m}_t \in \mathbb{R}^k$ and $\boldsymbol{K}_t \in \mathbb{R}^{k\times k}$.

The mean $\boldsymbol{m}_t$ of this distribution is defined by the mean function $m : \mathbb{R}^d \to \mathbb{R}$ of the Gaussian process:

$$\boldsymbol{m}_t = (m(\boldsymbol{t}_1), \ldots, m(\boldsymbol{t}_k))^\mathsf{T}.$$

Similarly, the covariance matrix $\boldsymbol{K}_t$ is defined by the covariance function $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$:

$$\boldsymbol{K}_t = \begin{pmatrix} k(\boldsymbol{t}_1, \boldsymbol{t}_1) & k(\boldsymbol{t}_1, \boldsymbol{t}_2) & \ldots & k(\boldsymbol{t}_1, \boldsymbol{t}_n) \\ k(\boldsymbol{t}_2, \boldsymbol{t}_1) & k(\boldsymbol{t}_2, \boldsymbol{t}_2) & \ldots & k(\boldsymbol{t}_2, \boldsymbol{t}_n) \\ \ldots & \ldots & \ldots & \ldots \\ k(\boldsymbol{t}_n, \boldsymbol{t}_1) & k(\boldsymbol{t}_n, \boldsymbol{t}_2) & \ldots & k(\boldsymbol{t}_n, \boldsymbol{t}_n) \end{pmatrix}. \tag{1}$$

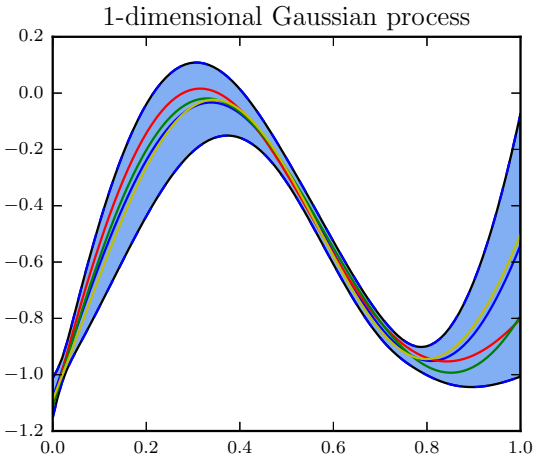

Figure 1 One-dimensional Gaussian processes

It's straightforward then, that a Gaussian process is completely defined by its mean and covariance functions. We will use the following notation:

$$f \sim \mathcal{GP}(m(\cdot), k(\cdot, \cdot))$$

While the mean function m can be an arbitrary real-valued function, the covariance function k has to be a kernel, so that the covariance matrices (1) it implies are symmetric and positive definite.

Fig. 1 shows an example of a one-dimensional Gaussian process. The dark blue line is the mean function of the process, the light blue region is the 3σ-region, and different color curves are samples from the process.

2.2 Gaussian process classification

Now we apply Gaussian processes to a binary classification problem. Suppose, we have a dataset $\{(\boldsymbol{x}_i, y_i) \,|\, i = 1, \ldots, n\}$, where $\boldsymbol{x}_i \in \mathbb{R}^d$, $y_i \in \{-1, 1\}$. Denote the matrix comprised of points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ by $\boldsymbol{X} \in \mathbb{R}^{n \times d}$ and the vector of corresponding class labels $y_1, \ldots, y_n$ by $\boldsymbol{y} \in \{-1, 1\}^n$. The task is to predict the class label $y_* \in \{-1, 1\}$ at a new point $\boldsymbol{x}_* \in \mathbb{R}^d$.

We consider the following model. First we introduce a latent function $f : \mathbb{R}^d \to \mathbb{R}$ and put a zero-mean GP prior over it:

$$f \sim \mathcal{GP}(0, k(\cdot, \cdot))$$

for some covariance function $k(\cdot, \cdot)$. For now the covariance function is supposed to be fixed.

Then we consider the probability of the object $\boldsymbol{x}_*$ belonging to positive class to be equal to $\sigma(f(\boldsymbol{x}_*))$ for the chosen sigmoid function σ:

$$p(y_* = +1 \,|\, \boldsymbol{x}_*) = \sigma(f(\boldsymbol{x}_*)). \tag{2}$$

In this work we use the logistic function $\sigma(z) = (1 + \exp(-z))^{-1}$, however one could use other sigmoid functions as well.

The probabilistic model for this setting is given by

$$p(\boldsymbol{y}, \boldsymbol{f} \,|\, \boldsymbol{X}) = p(\boldsymbol{y} \,|\, \boldsymbol{f}) p(\boldsymbol{f} \,|\, \boldsymbol{X}) = p(\boldsymbol{f} \,|\, \boldsymbol{X}) \prod_{i=1}^{n} p(y_i \,|\, f_i), \tag{3}$$

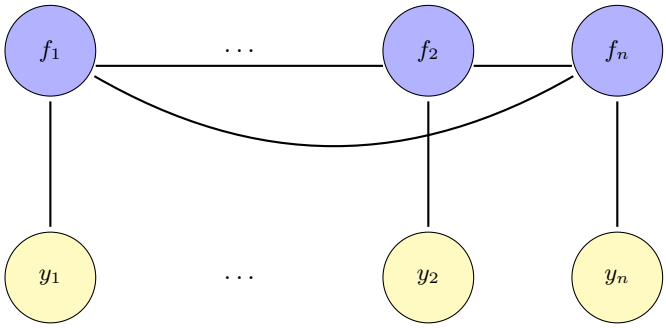

Figure 2 GP-classification graphical model

where $p(y_i \mid f_i)$ is sigmoid likelihood (2) and $p(\boldsymbol{f} \mid \boldsymbol{X}) = \mathcal{N}(\boldsymbol{f} \mid \boldsymbol{0}, K(\boldsymbol{X}, \boldsymbol{X}))$ is GP prior. The corresponding probabilistic graphical model is given in fig. 2.

Now inference in the model (3) can be done in two steps. First, for new data point $\boldsymbol{x}_*$ we should find the conditional distribution of the corresponding value of the latent process f_*. This can be done as follows

$$p(f_* \mid \boldsymbol{y}, \boldsymbol{X}, \boldsymbol{x}_*) = \int p(f_* \mid \boldsymbol{f}, \boldsymbol{X}, \boldsymbol{x}_*) p(\boldsymbol{f} \mid \boldsymbol{y}, \boldsymbol{X}) d\boldsymbol{f}. \tag{4}$$

Second, the probability that $\boldsymbol{x}_*$ belongs to the positive class is obtained by marginalizing over the latent variable f_*.

$$p(y_* = +1 \mid \boldsymbol{y}, \boldsymbol{X}, \boldsymbol{x}_*) = \int \sigma(f_*) p(f_* \mid \boldsymbol{y}, \boldsymbol{X}, \boldsymbol{x}_*) df_*. \tag{5}$$

Unfortunately, both integrals (4) and (5) are intractable since they involve a product of sigmoid functions and normal distributions. Thus, we have to use some integral-approximation techniques to estimate the predictive distribution.

For example, one can use Laplace approximation method, which builds a Gaussian approximation $q(\boldsymbol{f} \mid \boldsymbol{y}, \boldsymbol{X})$ to the true posterior $p(\boldsymbol{f} \mid \boldsymbol{y}, \boldsymbol{X})$. Substituting this Gaussian approximation back into (4) we obtain a tractable integral. The predictive distribution (5) remains intractable, but since this is a one-dimensional integral it can be easily estimated by quadratures or other techniques. The more detailed derivation of this algorithm and another algorithm, based on Expectation Propagation, can be found in [5].

Computational complexity of computing the predictive distribution both for the Laplace approximation method and Expectation propagation scales as $\mathcal{O}(n^3)$ since they both require to invert $n \times n$ matrix $K(\boldsymbol{X}, \boldsymbol{X})$. In section 3 we describe the concept of inducing points aimed to reduce this complexity.

2.3 Model adaptation

In the previous section we described how to fit a Gaussian process to the data in the classification problem. However, we only considered Gaussian processes with fixed covariance functions. This model can be rather limiting.

Most of the popular covariance functions have a set of parameters, which we refer to as covariance (or kernel) hyper-parameters. For example, the squared exponential covariance function

$$k_{SE}(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\theta}) = \sigma^2 \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{x}'\|^2}{l^2}\right)$$

has two parameters $\boldsymbol{\theta}$ – variance σ and length-scale l. An example of a more complicated popular covariance function is the Matern function, given by

$$k_{Matern}(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\theta}) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}\|\boldsymbol{x} - \boldsymbol{x}'\|}{l} \right)^{\nu} K_{\nu} \left(\frac{\sqrt{2\nu}}{\|\boldsymbol{x} - \boldsymbol{x}'\|} l \right),$$

with two positive parameters $\boldsymbol{\theta} = (\nu, l)$. Here K_ν is a modified Bessel function.

In order to get a good model for the data, one should find a good set of kernel hyper-parameters $\boldsymbol{\theta}$. Bayesian paradigm provides a way of tuning the kernel hyper-parameters of the GP-model through maximization of the model evidence (marginal likelihood), that is given by

$$p(\boldsymbol{y} \mid \boldsymbol{X}, \boldsymbol{\theta}) = \int p(\boldsymbol{y} \mid \boldsymbol{f}) p(\boldsymbol{f} \mid \boldsymbol{X}, \boldsymbol{\theta}) d\boldsymbol{f} \to \max_{\boldsymbol{\theta}}. \tag{6}$$

However, this integral is intractable for the model (3) since it involves a product of sigmoid functions and normal distribution. In subsequent sections we describe several methods to construct a variational lower bound to the marginal likelihood. Maximizing this lower bound with respect to kernel hyper-parameters $\boldsymbol{\theta}$, one could fit the model to the data.

3 Variational inducing point GP-classification

In the previous section we showed how Gaussian processes can be applied to solve classification problems. The computational complexity of GP for classification scales as $\mathcal{O}(n^3)$, that makes this method inapplicable to big data problems.

A number of approximate methods have been proposed in the literature for both GP-regression and GP-classification [8–10]. In this paper we consider methods based on the concept of inducing inputs. These methods construct an approximation based on the values of the process at some $m < n$ points. These points are referred to as inducing points. The idea is the following. The hidden Gaussian process f corresponds to some smooth low-dimensional surface in $\mathbb{R}^d$. This surface can in fact be well approximated by another Gaussian process with properly chosen m training points $\boldsymbol{Z} = (\boldsymbol{z}_1, \ldots, \boldsymbol{z}_m)^\intercal \in \mathbb{R}^{m \times d}$ and process values at that points $\boldsymbol{u} = (u_1, \ldots, u_m)^T$ (inducing inputs). Then predictions of this new process at training points are used for constructing approximate posterior distribution for $p(\boldsymbol{f} \mid \boldsymbol{y}, \boldsymbol{X})$. The positions $\boldsymbol{Z}$ of inducing inputs can be learned within training procedure. However, for simplicity in the following we clusterize the dataset $\boldsymbol{X}$ into m clusters using K-means and choose $\boldsymbol{Z}$ to be cluster centres. In practice we observe that this approach works well in almost all the cases.

3.1 Evidence lower bound

In the following we use a variational approach for solving maximum evidence problem (6). In this approach an evidence lower bound is introduced, that is simpler to compute than the evidence itself. Then this lower bound is maximized w.r.t. kernel hyperparameters $\boldsymbol{\theta}$ and additional variational parameters used for constructing the lower bound.

Let's consider the following augmented probabilistic model:

$$p(\boldsymbol{y}, \boldsymbol{f}, \boldsymbol{u} \mid \boldsymbol{X}, \boldsymbol{Z}) = p(\boldsymbol{y} \mid \boldsymbol{f}) p(\boldsymbol{f}, \boldsymbol{u} \mid \boldsymbol{X}, \boldsymbol{Z}) = \prod_{i=1}^{n} p(y_i \mid f_i) p(\boldsymbol{f}, \boldsymbol{u} \mid \boldsymbol{X}, \boldsymbol{Z}). \tag{7}$$

The graphical model for the model (7) is shown in fig. 3. Note that marginalizing the model (7) w.r.t $\boldsymbol{u}$ gives the initial model (3).

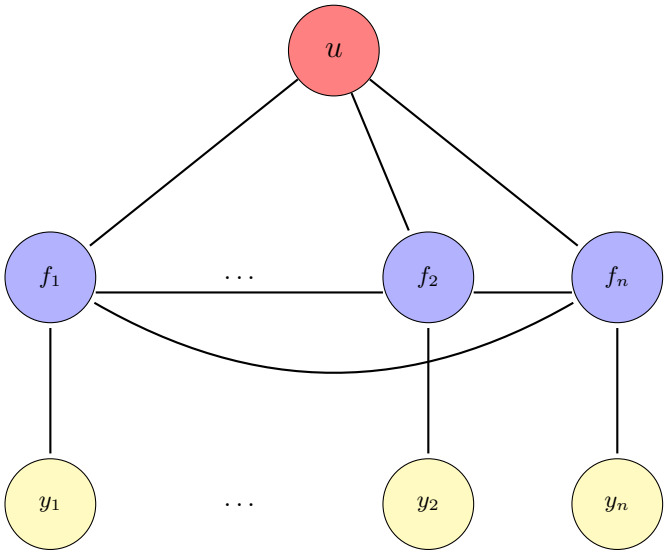

Figure 3 GP-classification graphical model

We denote the covariance matrix comprised of pairwise values of the covariance function $k(\cdot,\ \cdot)$ on the points $\boldsymbol{Z}$ by $K(\boldsymbol{Z}, \boldsymbol{Z}) = \boldsymbol{K}_{mm} \in \mathbb{R}^{m\times m}$. Similarly, we define $\boldsymbol{K}_{nn} = K(\boldsymbol{X}, \boldsymbol{X}) \in \mathbb{R}^{n\times n}$ and $\boldsymbol{K}_{nm} = K(\boldsymbol{X}, \boldsymbol{Z}) = \boldsymbol{K}_{mn}^{\mathsf{T}} \in \mathbb{R}^{n\times m}$.

As $\boldsymbol{u}$ and $\boldsymbol{f}$ are generated from the same Gaussian process with zero-mean prior,

$$\begin{aligned}
&p(\boldsymbol{f}, \boldsymbol{u} \,|\, \boldsymbol{X}, \boldsymbol{Z}) = \mathcal{N}([\boldsymbol{f}, \boldsymbol{u}] \,|\, [\boldsymbol{0}, \boldsymbol{0}], K([\boldsymbol{X}, \boldsymbol{Z}], [\boldsymbol{X}, \boldsymbol{Z}])),\\
&p(\boldsymbol{u} \,|\, \boldsymbol{Z}) = \mathcal{N}(\boldsymbol{u} \,|\, \boldsymbol{0}, \boldsymbol{K}_{mm}),\\
&p(\boldsymbol{f} \,|\, \boldsymbol{u}, \boldsymbol{X}, \boldsymbol{Z}) = \mathcal{N}(\boldsymbol{f} \,|\, \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{u}, \tilde{\boldsymbol{K}}),
\end{aligned} \tag{8}$$

where $\tilde{\boldsymbol{K}} = \boldsymbol{K}_{nn} - \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn}$. In the following for simplicity we omit dependence on $\boldsymbol{X}$ and $\boldsymbol{Z}$ in all formulas. Note that we do not consider optimization w.r.t. these values.

Applying the standard variational lower bound (see, for example [3]) to the augmented model (7), we obtain the following inequality:

$$\log p(\boldsymbol{y}) \geqslant \mathsf{E}_{q(\boldsymbol{u},\boldsymbol{f})} \log \frac{p(\boldsymbol{y}, \boldsymbol{u}, \boldsymbol{f})}{q(\boldsymbol{u}, \boldsymbol{f})} = \mathsf{E}_{q(\boldsymbol{u},\boldsymbol{f})} \log p(\boldsymbol{y} \,|\, \boldsymbol{f}) - \mathrm{KL}\left(q(\boldsymbol{u}, \boldsymbol{f}) \| p(\boldsymbol{u}, \boldsymbol{f})\right)$$

for any distribution $q(\boldsymbol{u}, \boldsymbol{f})$. This inequality becomes equality for the true posterior distribution $q(\boldsymbol{u}, \boldsymbol{f}) = p(\boldsymbol{u}, \boldsymbol{f} \,|\, \boldsymbol{y})$. Next we restrict the variational distribution $q(\boldsymbol{u}, \boldsymbol{f})$ to be of the form

$$q(\boldsymbol{u}, \boldsymbol{f}) = p(\boldsymbol{f} \,|\, \boldsymbol{u})q(\boldsymbol{u}), \tag{9}$$

where $q(\boldsymbol{u}) = \mathcal{N}(\boldsymbol{u} \,|\, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ for some $\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m\times m}$ and $p(\boldsymbol{f} \,|\, \boldsymbol{u})$ is determined by (8). This is the key approximation step in inducing points approach for Gaussian processes. The chosen family (9) subsumes that with large enough m all information about the hidden process values $\boldsymbol{f}$ at training points can be successfully restored from the values $\boldsymbol{u}$ at inducing inputs, i.e. $p(\boldsymbol{f} \,|\, \boldsymbol{u}, \boldsymbol{y}) \approx p(\boldsymbol{f} \,|\, \boldsymbol{u})$.

The form (9) of the variational distribution implies a Gaussian marginal distribution

$$q(\boldsymbol{f}) = \int p(\boldsymbol{f} \,|\, \boldsymbol{u})q(\boldsymbol{u})d\boldsymbol{u} = \mathcal{N}(\boldsymbol{f} \,|\, \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{\mu}, \boldsymbol{K}_{nn} + \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}(\boldsymbol{\Sigma} - \boldsymbol{K}_{mm})\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn}).$$

As $\log p(\boldsymbol{y} \mid \boldsymbol{f})$ depends on $\boldsymbol{u}$ only through $\boldsymbol{f}$, the expectation

$$\mathsf{E}_{q(\boldsymbol{u},\boldsymbol{f})} \log p(\boldsymbol{y} \mid \boldsymbol{f}) = \mathsf{E}_{q(\boldsymbol{f})} \log p(\boldsymbol{y} \mid \boldsymbol{f}) = \sum_{i=1}^{n} \mathsf{E}_{q(f_i)} \log p(y_i \mid f_i),$$

where $q(f_i)$ is the marginal distribution of $q(\boldsymbol{f})$:

$$q(f_i) = \mathcal{N}(f_i \mid \boldsymbol{k}_i^{\mathsf{T}} \boldsymbol{K}_{mm}^{-1} \boldsymbol{\mu}, \boldsymbol{K}_{ii} + \boldsymbol{k}_i^T \boldsymbol{K}_{mm}^{-1} (\boldsymbol{\Sigma} - \boldsymbol{K}_{mm}) \boldsymbol{K}_{mm}^{-1} \boldsymbol{k}_i) = \mathcal{N}(f_i \mid m_i, S_i^2), \tag{10}$$

and $\boldsymbol{k}_i$ is the i-th column of matrix $\boldsymbol{K}_{mn}$.

Finally,

$$\mathrm{KL}\left(q(\boldsymbol{u}, \boldsymbol{f}) \| p(\boldsymbol{u}, \boldsymbol{f})\right) = \mathrm{KL}\left(q(\boldsymbol{u}) p(\boldsymbol{f} \mid \boldsymbol{u}) \| p(\boldsymbol{u}) p(\boldsymbol{f} \mid \boldsymbol{u})\right) = \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right).$$

Combining everything back together, we obtain the evidence lower bound

$$\log p(\boldsymbol{y}) \geqslant \sum_{i=1}^{n} \mathsf{E}_{q(f_i)} \log p(y_i \mid f_i) - \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right). \tag{11}$$

Note that the KL-divergence term in the lower bound (11) can be computed analytically since it is a KL-divergence between two normal distributions. In order to compute the expectations $\mathsf{E}_{q(f_i)} \log p(y_i \mid f_i)$, we have to use integral approximating techniques.

The evidence lower bound (ELBO) (11) can be maximized with respect to variational parameters $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ and kernel hyper-parameters. Using the optimal distribution $q(\boldsymbol{u})$, we can perform predictions for new data point $\boldsymbol{x}_*$ as follows

$$p(f_* \mid \boldsymbol{y}) = \int p(f_* \mid \boldsymbol{u}, \boldsymbol{f}) p(\boldsymbol{u}, \boldsymbol{f} \mid \boldsymbol{y}) d\boldsymbol{u} d\boldsymbol{f} \approx \int p(f_* \mid \boldsymbol{u}, \boldsymbol{f}) q(\boldsymbol{u}, \boldsymbol{f}) d\boldsymbol{u} d\boldsymbol{f} =$$
$$= \int p(f_* \mid \boldsymbol{u}, \boldsymbol{f}) p(\boldsymbol{f} \mid \boldsymbol{u}) q(\boldsymbol{u}) d\boldsymbol{u} d\boldsymbol{f} = \int p(f_* \mid \boldsymbol{u}) q(\boldsymbol{u}) d\boldsymbol{u}.$$

The last integral is tractable since both terms $p(f_* \mid \boldsymbol{u})$ and $q(\boldsymbol{u})$ are normal distributions.

Note, that in case of regression with Gaussian noise the distributions $p(y_i \mid f_i)$ are Gaussians and thus the expectations $\mathsf{E}_{q(f_i)} \log p(y_i \mid f_i)$ are tractable. Paper [1] suggests maximization of the lower bound (11) with respect to $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ and covariance hyper-parameters with stochastic optimization techniques for GP-regression.

3.2 SVI method

In case of classification we can't analytically compute the expectations $\mathsf{E}_{q(f_i)} \log p(y_i \mid f_i)$ in the lower bound (11). However, the expectations are one-dimensional Gaussian integrals and can thus be effectively approximated with a range of techniques. In paper [2] Gauss-Hermite quadratures are used for this purpose. Note that the lower bound (11) has the form "sum over training objects". Hence this bound can be maximized using stochastic optimization techniques. Paper [2] suggests to maximize the lower bound (11) with respect to the variational parameters $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ and kernel hyper-parameters $\boldsymbol{\theta}$ using stochastic optimization. We refer to this method as `svi` method (abbreviation for Stochastic Variational Inference). The lower bound (11) and all it's derivatives can be computed in $\mathcal{O}(nm^2 + m^3)$. This complexity has a linear dependence on n, hence `svi` method can be applied for the case of big training data.

4 Tractable evidence lower bound for GP-classification

In the previous section we described the `svi` method. It is based on stochastic optimization of the lower bound (11) for marginal likelihood and the lower bound itself is computed in $\mathcal{O}(nm^2)$. But the bound depends on $\mathcal{O}(m^2)$ parameters which makes the optimization problem hard to solve when a big number of inducing points is needed.

For GP-regression the situation is similar. Paper [1] describes a method analogical to the `svi` method for classification. The only difference is that the lower bound becomes tractable in case of regression. Then the paper [6] tries to solve the problem of big $\mathcal{O}(m^2)$ number of parameters in the algorithm from [1] in the following way. In case of regression the lower bound (11) can be analytically optimised with respect to variational parameters $\boldsymbol{\mu}, \boldsymbol{\Sigma}$. Doing so and substituting the optimal values back into the lower bound, one can obtain a new lower bound to the marginal likelihood that depends solely on kernel hyper-parameters $\boldsymbol{\theta}$. This simplifies the optimization problem by dramatically reducing the number of optimization parameters. Unfortunately, this new bound doesn't have a form of "sum over objects", hence stochastic optimization methods are no longer applicable here. However, in our experiments we've found that even for fairly big datasets the method from [6] outperforms [1] despite the lack of stochastic optimization.

In the following subsection we devise an approach similar to the method of [6] for the case of classification. We provide a tractable evidence lower bound and analytically maximize it with respect to variational parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Substituting the optimal values of these parameters back into the lower bound we obtain a new lower bound, that depends only on kernel hyper-parameters $\boldsymbol{\theta}$.

4.1 Global evidence lower bound

In order to derive a tractable lower bound for (11), we seek a quadratic approximation to the log-logistic function $\log p(y_i \,|\, f_i) = \log \sigma(y_i f_i)$, where $\log \sigma(t) = -\log(1 + \exp(-t))$. The paper [4] provides a global parametric quadratic lower bound for this function:

$$\log \sigma(t) \geqslant \frac{t}{2} - \frac{\xi_t}{2} + \log \sigma(\xi_t) - \lambda(\xi_t)(t^2 - \xi_t^2), \ \forall t$$

where $\lambda(\xi_t) = \tanh(\xi_t)/(4\xi_t)$ and $\xi_t \in \mathbb{R}$ is a parameter of the bound. This bound is tight when $t^2 = \xi_t^2$.

Substituting this bound back to (11) with separate values $\boldsymbol{\xi} = \{\xi_i \,|\, i = 1, \ldots, n\}$ for every data point, we obtain a tractable lower bound

$$\log p(y) \geqslant \sum_{i=1}^{n} \mathsf{E}_{q(f_i)} \log p(y_i|f_i) - \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right) = \sum_{i=1}^{n} \mathsf{E}_{q(f_i)} \log \sigma(y_i f_i) - \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right) \geqslant$$

$$\geqslant \sum_{i=1}^{n} \left(\mathsf{E}_{q(f_i)} \left[\log \sigma(\xi_i) + \frac{y_i f_i - \xi_i}{2} - \lambda(\xi_i)(f_i^2 - \xi_i^2) \right] \right) - \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right) =$$

$$= \sum_{i=1}^{n} \left(\log \sigma(\xi_i) - \frac{\xi_i}{2} + \lambda(\xi_i)\xi_i^2 \right) + \frac{1}{2}\boldsymbol{\mu}^{\mathrm{T}} \boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn} \boldsymbol{y} - \mathrm{tr}\left(\boldsymbol{\Lambda}(\boldsymbol{\xi})(\boldsymbol{K}_{nn} + \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}(\boldsymbol{\Sigma} - \boldsymbol{K}_{mm})\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn}) \right) -$$

$$- \boldsymbol{\mu}^{\mathrm{T}} \boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn} \boldsymbol{\Lambda}(\boldsymbol{\xi}) \boldsymbol{K}_{nm} \boldsymbol{K}_{mm}^{-1} \boldsymbol{\mu} - \mathrm{KL}\left(q(\boldsymbol{u}) \| p(\boldsymbol{u})\right) = J(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\xi}, \boldsymbol{\theta}),$$

where

$$\boldsymbol{\Lambda}(\boldsymbol{\xi}) = \begin{pmatrix} \lambda(\xi_1) & 0 & \ldots & 0 \\ 0 & \lambda(\xi_2) & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & \lambda(\xi_n) \end{pmatrix}.$$

Algorithm 1 vi-JJ method

Input: n_{upd}, n_{fun}
Output: $\boldsymbol{\theta}$, $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$
  $\boldsymbol{\mu}$, $\boldsymbol{\Sigma} \leftarrow \mathbf{0}$, $\boldsymbol{I}$
  repeat
    $\tilde{\boldsymbol{\mu}}$, $\tilde{\boldsymbol{\Sigma}} \leftarrow \boldsymbol{\mu}$, $\boldsymbol{\Sigma}$
    for $j \leftarrow 1, \ldots, n_{upd}$: *// stage 1: updating $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, $\boldsymbol{\xi}$*
      m_t, $S_t \leftarrow \boldsymbol{k}_t^{\intercal} \boldsymbol{K}_{mm}^{-1} \tilde{\boldsymbol{\mu}}$, $\boldsymbol{K}_{tt} + \boldsymbol{k}_t^T \boldsymbol{K}_{mm}^{-1} (\tilde{\boldsymbol{\Sigma}} - \boldsymbol{K}_{mm}) \boldsymbol{K}_{mm}^{-1} \boldsymbol{k}_t, \quad t = 1, \ldots, n$
      $\tilde{\xi}_t^{\,2} \leftarrow m_t^2 + S_t^2, \quad t = 1, \ldots, n$
    $\tilde{\boldsymbol{\mu}}$, $\tilde{\boldsymbol{\Sigma}} \leftarrow \hat{\boldsymbol{\mu}}(\tilde{\boldsymbol{\xi}})$, $\hat{\boldsymbol{\Sigma}}(\tilde{\boldsymbol{\xi}})$ *// see (12), (13).*
    $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, $\boldsymbol{\xi} \leftarrow \tilde{\boldsymbol{\mu}}$, $\tilde{\boldsymbol{\Sigma}}$, $\tilde{\boldsymbol{\xi}}$
    $\boldsymbol{\theta}$ = minimize($\hat{J}(\cdot, \boldsymbol{\xi})$, method='L-BFGS-B', maxfun=n_{fun}) *// stage 2: updating $\boldsymbol{\theta}$*
  until convergence

Differentiating J with respect to $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ and setting the derivatives to zero, we obtain

$$\hat{\boldsymbol{\Sigma}}(\boldsymbol{\xi}) = (2\boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn} \boldsymbol{\Lambda}(\boldsymbol{\xi}) \boldsymbol{K}_{nm} \boldsymbol{K}_{mm}^{-1} + \boldsymbol{K}_{mm}^{-1})^{-1}, \tag{12}$$

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\xi}) = \frac{1}{2} \hat{\boldsymbol{\Sigma}}(\boldsymbol{\xi}) \boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn} \boldsymbol{y}. \tag{13}$$

Substituting the optimal values of variational parameters back to the lower bound J and omitting the terms not depending on $\boldsymbol{\theta}$ and $\boldsymbol{\xi}$, we obtain a compact lower bound

$$\begin{aligned} \hat{J}(\boldsymbol{\theta}, \boldsymbol{\xi}) = \sum_{i=1}^{n} \left(\log \sigma(\xi_i) - \frac{\xi_i}{2} + \lambda(\xi_i) \xi_i^2 \right) + \frac{1}{8} \boldsymbol{y}^{\intercal} \boldsymbol{K}_{nm} \boldsymbol{B}^{-1} \boldsymbol{K}_{mn} \boldsymbol{y} \\ + \frac{1}{2} \log |\boldsymbol{K}_{mm}| - \frac{1}{2} \log |\boldsymbol{B}| - \operatorname{tr}(\boldsymbol{\Lambda}(\boldsymbol{\xi}) \tilde{\boldsymbol{K}}), \end{aligned} \tag{14}$$

where

$$\tilde{\boldsymbol{K}} = \boldsymbol{K}_{nn} - \boldsymbol{K}_{nm} \boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn},$$
$$\boldsymbol{B} = 2\boldsymbol{K}_{mn} \boldsymbol{\Lambda}(\boldsymbol{\xi}) \boldsymbol{K}_{nm} + \boldsymbol{K}_{mm}.$$

In the following we consider three different methods for maximizing the lower bound $\hat{J}(\boldsymbol{\theta}, \boldsymbol{\xi})$.

Note, that given the values of $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, $\boldsymbol{\theta}$, we can maximize $J(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\xi}, \boldsymbol{\theta})$ with respect to $\boldsymbol{\xi}$ analytically. The optimal values for $\boldsymbol{\xi}$ are given by

$$\xi_i^2 = \mathsf{E}_{q(\boldsymbol{f})} f_i^2 = m_i^2 + S_i^2. \tag{15}$$

The values m_i and S_i were defined in (10). In our first method we use analytical formulas (15) to recompute the values of $\boldsymbol{\xi}$ and use gradient-based optimization to maximize the bound with respect to $\boldsymbol{\theta}$. The pseudocode is given in Alg. 1. We will refer to this method as `vi-JJ`, where `JJ` stands for Jaakkola and Jordan, the authors of [4]. Note that the computational complexity of one iteration of this method is $\mathcal{O}(nm^2)$, the same as for the `svi` method.

Our second method uses gradient-based optimization to maximize $\hat{J}$ with respect to both $\boldsymbol{\theta}$ and $\boldsymbol{\xi}$. Note that in this method we don't have to recompute $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ at each iteration which makes the methods iterations empirically faster for big values of m. We refer to this method as `vi-JJ-full`.

Finally, `vi-JJ-hybrid` is a combination of the two methods described above. The general scheme of this method is the same as `vi-JJ`. In the `vi-JJ-hybrid` method we use analytical formulas to recompute $\boldsymbol{\xi}$ as we do in the `vi-JJ` method at stage 1, but at stage 2 we use gradient-based optimization with respect to both $\boldsymbol{\xi}$ and $\boldsymbol{\theta}$. The virtues of this method will be described in the experiments section.

4.2 Tractable local approximation to the evidence lower bound

Another way to obtain a tractable approximation to the lower bound (11) is to use a local quadratic approximation for the log-logistic function $\log p(y_i \,|\, f_i)$. In this way we perform a second-order Taylor expansion of this function at points $\boldsymbol{\xi} = \{\xi_i \,|\, i = 1, \ldots, n\}$:

$$\log p(y_i \,|\, f_i) \approx -\log(1+\exp(-y_i\xi_i)) + \frac{y_i}{1+\exp(y_i\xi_i)}(f_i - \xi_i) - \frac{y_i^2 \exp(y_i\xi_i)}{2(1+\exp(y_i\xi_i))^2}(f_i-\xi_i)^2. \quad (16)$$

The following derivation is analogical to the derivation in the previous section. Substituting the approximation (16) into the lower bound (11), we obtain

$$\log p(y) \geqslant \sum_{i=1}^{n} \mathsf{E}_{q(f_i)} \log p(y_i|f_i) - \mathrm{KL}\left(q(\boldsymbol{u})\|p(\boldsymbol{u})\right) \approx$$

$$\approx -\sum_{i=1}^{n} \log(1+\exp(-y_i\xi_i)) + \varphi(\boldsymbol{\xi})^T(\boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{\mu} - \boldsymbol{\xi}) -$$

$$-\operatorname{tr}\left(\boldsymbol{\Psi}(\boldsymbol{\xi})(\boldsymbol{K}_{nn} + \boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}(\boldsymbol{\Sigma} - \boldsymbol{K}_{mm})\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn})\right) - (\boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{\mu} - \boldsymbol{\xi})^T\boldsymbol{\Psi}(\boldsymbol{\xi})(\boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1}\boldsymbol{\mu} - \boldsymbol{\xi}) -$$

$$-\frac{1}{2}\left(\log\frac{|\boldsymbol{K}_{mm}|}{|\boldsymbol{\Sigma}|} - m + \operatorname{tr}(\boldsymbol{K}_{mm}^{-1}\boldsymbol{\Sigma}) + \boldsymbol{\mu}^T\boldsymbol{K}_{mm}^{-1}\boldsymbol{\mu}\right).$$

Here $\boldsymbol{\Psi}(\boldsymbol{\xi})$ is a diagonal matrix

$$\boldsymbol{\Psi}(\boldsymbol{\xi}) = \begin{pmatrix} \psi(\xi_1) & 0 & \ldots & 0 \\ 0 & \psi(\xi_2) & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & \psi(\xi_n) \end{pmatrix},$$

where $\psi(\xi_i) = \dfrac{y_i^2 \exp(y_i\xi_i)}{2(1+\exp(y_i\xi_i))^2}$.

Differentiating the approximate bound with respect to $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, and $\boldsymbol{\xi}$, and setting the derivatives to zero, we obtain the following formulas for optimal values of these parameters:

$$\hat{\boldsymbol{\Sigma}}(\boldsymbol{\xi}) = \left(2\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn}\boldsymbol{\Psi}(\boldsymbol{\xi})\boldsymbol{K}_{nm}\boldsymbol{K}_{mm}^{-1} + \boldsymbol{K}_{mm}^{-1}\right)^{-1}, \quad (17)$$

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\xi}) = \hat{\boldsymbol{\Sigma}}(\boldsymbol{\xi})\boldsymbol{K}_{mm}^{-1}\boldsymbol{K}_{mn}\boldsymbol{v}(\boldsymbol{\xi}), \quad (18)$$

$$\xi_i = m_i.$$

Here

$$\boldsymbol{v}(\boldsymbol{\xi}) = \boldsymbol{\varphi}(\boldsymbol{\xi}) + 2\boldsymbol{\Psi}(\boldsymbol{\xi})\boldsymbol{\xi},$$

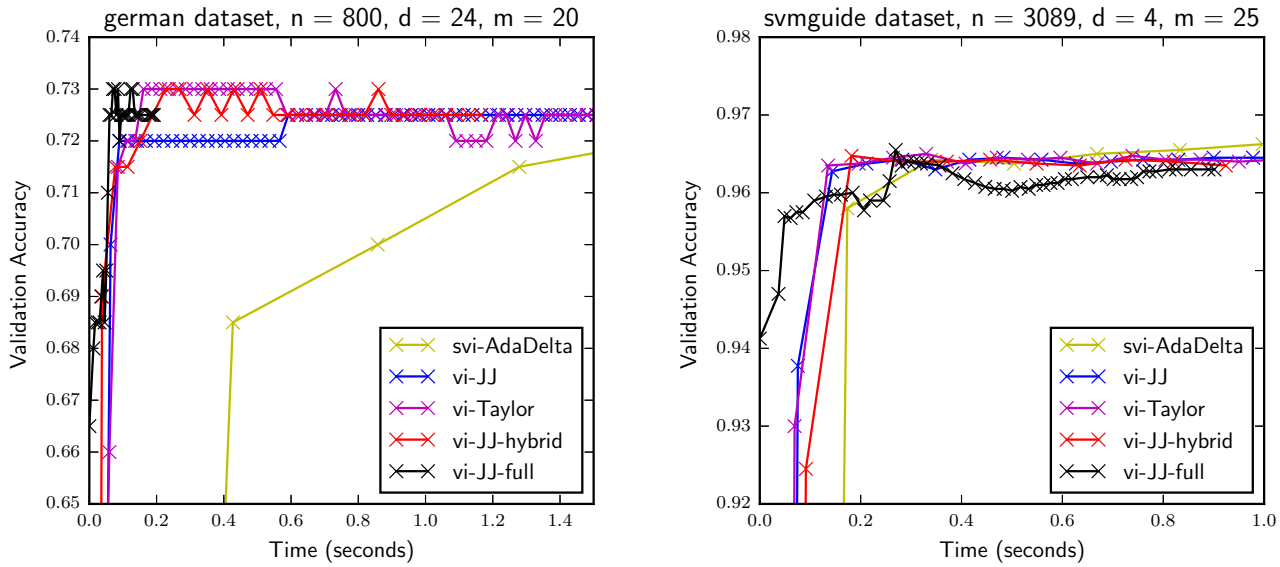

Figure 4 Methods performance on small datasets

and $\boldsymbol{\varphi}(\boldsymbol{\xi})$ is a vector, composed of

$$\boldsymbol{\varphi}(\boldsymbol{\xi})_i = \frac{y_i}{1 + \exp(y_i \xi_i)}.$$

Substituting the optimal values for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, given by (17), (18), back into the approximate bound, and omiting the terms, that do not depend on $\boldsymbol{\theta}$, we obtain the following approximate lower bound:

$$\tilde{J}_{\boldsymbol{\xi}} = \frac{1}{2}\boldsymbol{v}(\boldsymbol{\xi})^{\mathrm{T}}\boldsymbol{K}_{nm}\boldsymbol{B}^{-1}\boldsymbol{K}_{mn}\boldsymbol{v}(\boldsymbol{\xi}) + \frac{1}{2}\log|\boldsymbol{K}_{mm}| - \frac{1}{2}\log|\boldsymbol{B}| - \mathrm{tr}(\boldsymbol{\Psi}(\boldsymbol{\xi})\tilde{\boldsymbol{K}}), \tag{19}$$

where

$$\boldsymbol{B} = 2\boldsymbol{K}_{mn}\boldsymbol{\Psi}(\boldsymbol{\xi})\boldsymbol{K}_{nm} + \boldsymbol{K}_{mm}.$$

Note that the lower bound (19) is not a global lower bound for the log-evidence $\log p(y)$. However, locally we get a good approximation of the evidence lower bound (11).

For maximizing the approximate lower bound (19) we consider a method, analogical to `vi-JJ`. In order to specify this method we simply substitute the bound $\hat{J}(\cdot, \boldsymbol{\xi})$ by $\tilde{J}_{\boldsymbol{\xi}}$ in the second stage in Alg. 1. We will refer to this method as `vi-Taylor`. The computational complexity of one iteration of this method is once again $\mathcal{O}(nm^2)$.

5 Experiments

In this section we empirically compare the derived `vi-JJ`, `vi-Taylor`, `vi-JJ-full` and `vi-JJ-hybrid` methods with `svi`. Below we describe the setting of the experiments and discuss their results.

5.1 Experimental setting

In our experiments we compared 5 methods for variational inducing point GP-classification:

– `svi-AdaDelta` uses the AdaDelta optimization method for maximization of the lower bound (11) as it is done in the paper [2];

– `vi-JJ` was described in section 4.1;
– `vi-Taylor` was described in section 4.2;
– `vi-JJ-full` was described in section 4.1;
– `vi-JJ-hybrid` was described in section 4.1.

We also tried using deterministic L-BFGS-B optimization method for maximizing the evidence lower bound (11), but it worked substantially worse than all the other methods. Note, that all the methods have the same complexity of epochs $\mathcal{O}(nm^2)$. Table 1 shows which variables are optimized numerically and which are optimized analytically for each method.

In our experiments we didn't optimize the lower bound with respect to the positions $\boldsymbol{Z}$ of inducing points. Instead we used K-means clustering procedure with K equal to the number m of inducing inputs and took clusters centres as $\boldsymbol{Z}$. Also we used the squared exponential covariance function (see section 1) in all experiments with a Gaussian noise term.

The stochastic method `svi-AdaDelta` requires the user to manually specify the learning rate and the batch size for the optimization method. For the former we had to run the method with different learning rates and choose the value that resulted in the fastest convergence. We used the learning rates from a fixed grid with a step of 0.1. It always happened, that for the largest value from the grid the method diverged, and for the smallest the method converged slower, than for some medium value, verifying, that the optimal learning rate was somewhere in the range. To choose the batch size we used the following convention. For small *german* and *svmguide* datasets we've set the batch size to 50. For other datasets we used approximately $\frac{n}{100}$ as the batch size, where n is the size of the training set.

For the `vi-JJ`, `vi-Taylor` and `vi-JJ-hybrid` in all of the experiments on every iteration we recomputed the values of $\boldsymbol{\xi}$, $\boldsymbol{\mu}$, and $\boldsymbol{\Sigma}$ 3 times ($n_{upd} = 3$ in algorithm 1). To tune $\boldsymbol{\theta}$, on every iteration we've run L-BFGS-B optimization method constrained to do no more than 5 evaluations of the lower bound and it's gradient. We found that these values of parameters work well for all the datasets we experimented with.

For the `svi-AdaDelta` method we used optimization w.r.t. Cholesky factor of the matrix $\boldsymbol{\Sigma}$ in order to maintain it's positive definiteness, as described in [2]. We used AdaDelta optimization method implementation from the *climin* toolbox [7] as is done in the original paper.

For every dataset we experimented with a number of inducing points to verify that the results of the methods are close to optimal.

We evaluate the methods plotting the accuracy of their predictions on the test data against time. All of the plots have titles of the following format.

$$[\text{name of the dataset}], n = [\text{number of objects in the training set}],$$

$$d = [\text{number of features}], m = [\text{number of inducing inputs}]$$

Table 1 Methods outline

Method	Numerically optimized variables	Analytically optimized variables
`svi-AdaDelta`	$\boldsymbol{\theta}$, $\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$	
`vi-JJ`, `vi-Taylor`	$\boldsymbol{\theta}$	$\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$, $\boldsymbol{\xi} \in \mathbb{R}^n$
`vi-JJ-hybrid`	$\boldsymbol{\theta}$, $\boldsymbol{\xi} \in \mathbb{R}^n$	$\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$, $\boldsymbol{\xi} \in \mathbb{R}^n$
`vi-JJ-full`	$\boldsymbol{\theta}$, $\boldsymbol{\xi} \in \mathbb{R}^n$	$\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$

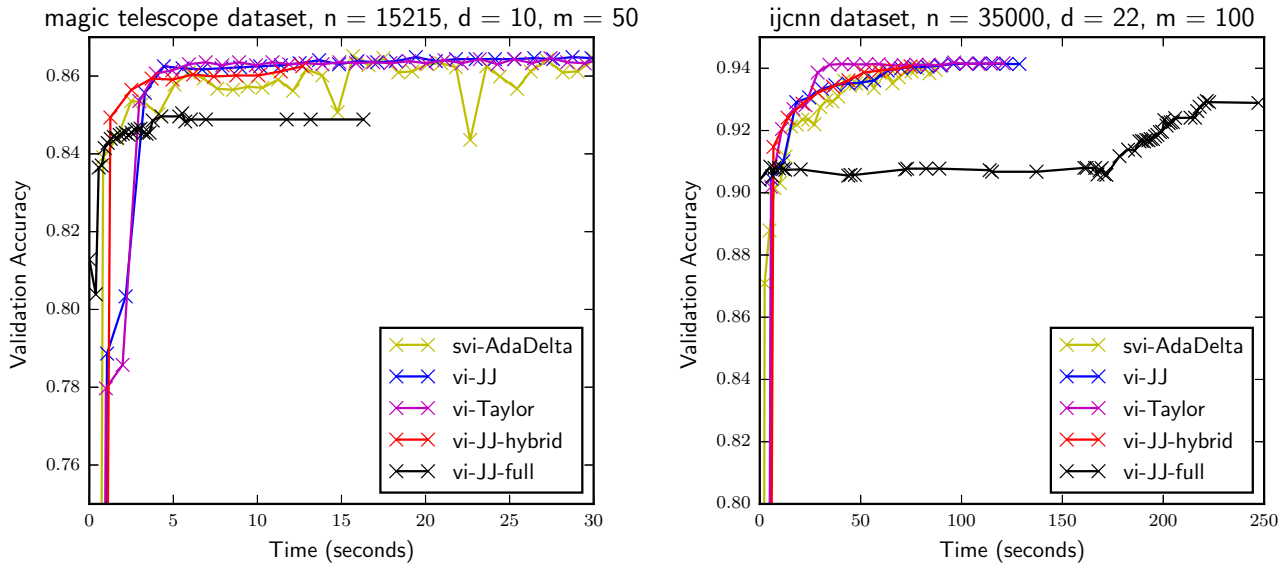

Figure 5 Methods performance on medium datasets

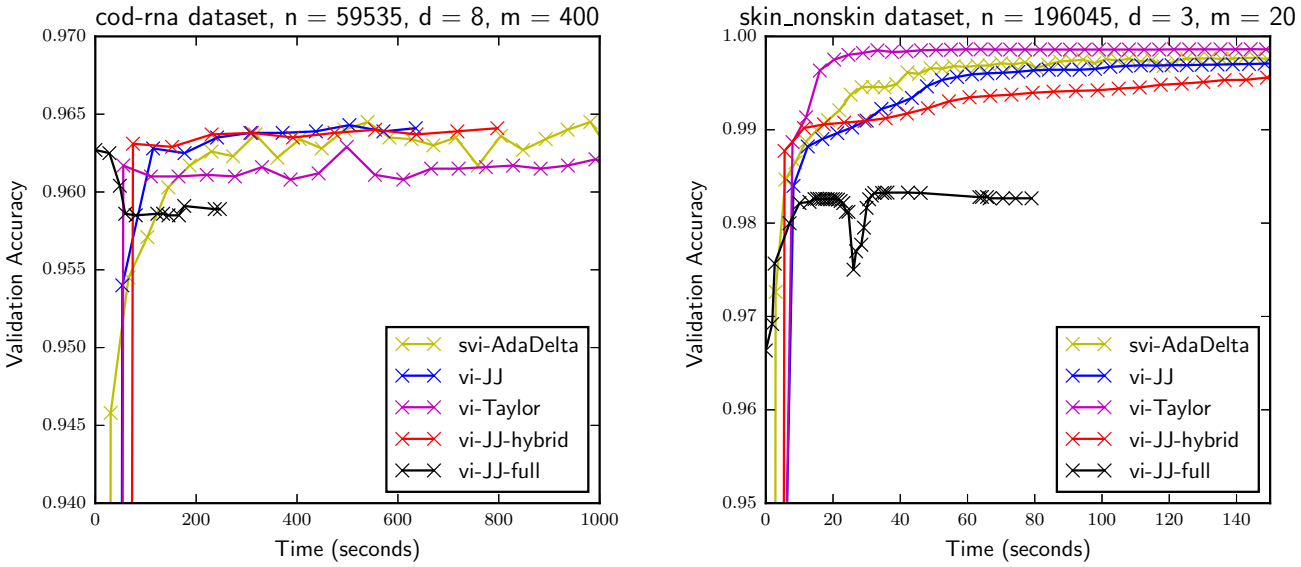

Figure 6 Methods performance on big datasets

We also preprocessed all the datasets by normalizing the features setting the mean of all features to 0 and the variance to 1. For datasets without available test data, we used 20% of the data as a test set and 80% as a train set.

5.2 Results and discussion

We compared the methods' performance on seven datasets. Here we discuss the results.

Fig. 4 provides the results for *german* and *svmguide* datasets. As we can see, on the small *german* dataset the stochastic `svi-AdaDelta` method struggles, and it takes it longer to achieve the optimal quality, than for all the other methods, which show similar results. On the *svmguide*

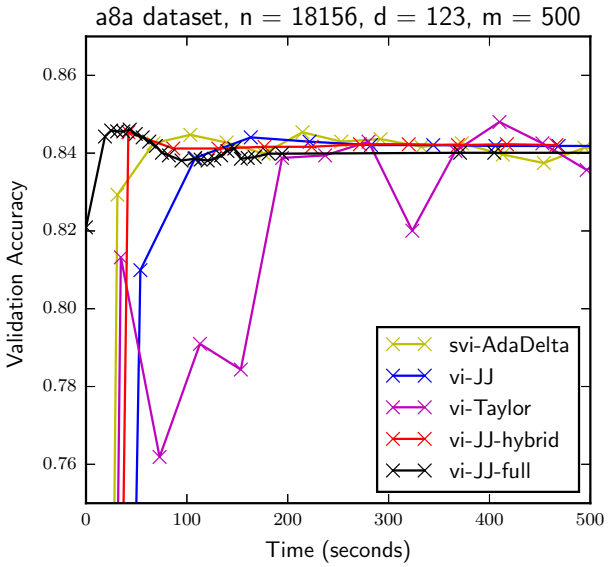

Figure 7 Methods performance on the a8a dataset

dataset it takes `vi-JJ-full` and `svi-AdaDelta` a little bit longer to converge, while the other three methods show roughly the same performance.

The results on *magic telescope* and *ijcnn* datasets are provided in fig. 5. On the *magic telescope* dataset the `vi-JJ` and `vi-Taylor` show poor quality on the first iterations, but still manage to converge faster, than `svi-AdaDelta`. On both datasets the `vi-JJ-hybrid` method works similar to `vi-JJ` and `vi-Taylor`, but shows better quality on first iterations on the *magic telescope* data. `vi-JJ-full` can't converge to a reasonable quality on both datasets.

Fig. 6 provides the results on big *cod-rna* and *skin_nonskin* datasets. On these datasets the `vi-JJ-full` once again fails to achieve a reasonable quality, while the other methods work similarly.

Finally, the results on *a8a* data are provided in fig. 7. Here we use a rather big amount of $m = 500$ inducing inputs. As we can see, the `vi-JJ-full` and `vi-JJ-hybrid` are the fastest to achieve the optimal quality. The `svi-AdaDelta` method also converges reasonably fast, while the `vi-JJ` and `vi-Taylor` struggle a little bit.

In general `vi-JJ`, `vi-Taylor` and `vi-JJ-hybrid` methods perform similar to `svi-AdaDelta` method. On the big dataset *skin_nonskin* with only 3 features the `vi-JJ-hybrid` is a little bit slower than the stochastic `svi-AdaDelta`, but on all the other datasets it is better. The `vi-Taylor` and `vi-JJ` struggle with *a8a*, but are otherwise comparable to `vi-JJ-hybrid`. The stochastic `svi-AdaDelta` method performs poorly on small datasets and even on the big *skin_nonskin* data doesn't manage to substantially outperform the other methods, even provided a good value of learning rate. Finally, `vi-JJ-full` works well on small data and on the *a8a*, but on all the other datasets it doesn't manage to achieve a reasonable quality.

6 Conclusion

In this paper we presented a new approach to training variational inducing input Gaussian process classification. We derived two new tractable evidence lower bounds and described several ways to maximize them. The resulting methods `vi-JJ`, `vi-JJ-full`, `vi-JJ-hybrid` and `vi-Taylor` are similar to the method of [6] for GP-regression.

We provided an experimental comparison of our methods with the current state-of-the-art method `svi-AdaDelta` of [2]. In our experimental setting, our approach proved to be more practical, as it converges to the optimal quality as fast as the `svi-AdaDelta` method without requiring the user to manually choose the parameters of the optimization method.

The four described `vi` methods showed similar performance and it's hard to distinguish them. However, note that the `vi-Taylor` approach is more general and can be applied to the likelihood functions, that are not logistic. We could also easily derive a method, similar to `vi-JJ-hybrid` and `vi-JJ-full` for the non-logistic case, but it is out of scope of this paper.

References

[1] J. Hensman, N. Fusi, and N. D. Lawrence. 2013. Gaussian processes for big data. *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*. P. 282–290.

[2] J. Hensman, G. Matthews, Z. Ghahramani. 2015. Scalable Variational Gaussian Process Classification. *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*.

[3] T. Jaakkola. 2000. Tutorial on Variational Approximation Methods. *Advanced Mean Field Methods: Theory and Practice*. P. 129–159.

[4] T. Jaakkola, M. Jordan. 1997. A Variational Approach to Bayesian Logistic Regression Models and Their Extensions. *Proceedings of the 1997 Conference on Artificial Intelligence and Statistics*.

[5] C. E. Rasmussen and C. K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA.

[6] M. Titsias. 2009. Variational learning of inducing variables in sparse Gaussian processes. *Proceedings of the Twelfth International Workshop on Artificial Intelligence and Statistics*. vol. 5, P. 567–574.

[7] Climin. URL: `http://github.com/BRML/climin`.

[8] A. Smola, P. Bartlett 2001. Sparse greedy Gaussian process regression. *Neural Information Processing Systems*.

[9] J. Quinonero-Candela, C. Rasmussen 2005. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*. vol. 6, P. 1939–1959.

[10] L. Csato, M. Opper 2002. Sparse online Gaussian processes. *Neural Computation*. vol. 14, P. 641–668.